

The Early Bird Gets the Worm: Major Retention in CS3

Fee Christoph
University of Michigan
Ann Arbor, Michigan, USA
felichri@umich.edu

Westley Weimer
University of Michigan
Ann Arbor, Michigan, USA
weimerw@umich.edu

Kevin Angstadt
St. Lawrence University
Canton, New York, USA
kangstadt@stlawu.edu

ABSTRACT

Enrollments continue to rise in Computer Science courses, yet fostering inclusive climates and retaining diverse student bodies remain key challenges. Gender ratios remain heavily skewed and many demographics are severely underrepresented. Numerous studies investigate student retention in introductory courses, but few focus on later stages of the CS retention “pipeline”, where techniques and findings from earlier courses may no longer apply.

In this work, we focus on the relatively understudied transition from introductory CS courses to upper level courses via CS3 (data structures and algorithms). We conduct an analysis of archival data for a CS3 course at a large, public university in the US, analyzing anonymized student assignments and university student records to identify factors that result in students choosing not to declare the major. Our analysis indicates that sex alone is not enough to predict students leaving the program after CS3 (despite reporting a desire to declare the major). However, we identify that students intending to major in CS who take CS3 later in their academic careers (often associated with non-traditional students) are 13% less likely to actually declare a CS major ($p = 0.00005$). Further, we find a disparity between these students and their “fast-tracked” counterparts in their project performance as measured by an autograder ($p = 0.00003$). Our findings indicate that the strategy of introducing students to CS early in their college careers and swiftly passing them through the intro sequence is effective in retaining students, yet unintentionally leaves behind those who reach CS in a more roundabout way.

CCS CONCEPTS

• **Social and professional topics** → **Computer science education**.

KEYWORDS

CS3; retention, fast-tracking; archival data

ACM Reference Format:

Fee Christoph, Westley Weimer, and Kevin Angstadt. 2021. The Early Bird Gets the Worm: Major Retention in CS3. In *26th ACM Conference on Innovation and Technology in Computer Science Education V. 1 (ITiCSE 2021)*, June 26–July 1, 2021, Virtual Event, Germany. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3430665.3456335>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ITiCSE 2021, June 26–July 1, 2021, Virtual Event, Germany

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8214-4/21/06...\$15.00

<https://doi.org/10.1145/3430665.3456335>

1 INTRODUCTION

While enrollment in undergraduate Computer Science courses continues to be at all-time highs and growing, retention of certain populations within the community remains a pressing issue [13, 14]. Studies show that marginalized populations in undergraduate Computer Science (e.g., women, underrepresented minorities, returning adult students, etc.) are more significantly affected by a low sense of belonging [24, 27, 34] and self efficacy [6, 7, 21, 27], and a perception of the major as asocial or unrelated to the real world [6, 11], leading to higher attrition rates among these populations [4, 34]. Supporting diverse student bodies is relevant not only for innovation and business [26] but also for retention and satisfaction [2, 4].

Previous research has primarily focused on retention in CS1 and CS2 [18, 36, 37]. Retention in higher-level courses, including CS3, appears as parts of longitudinal studies [2, 32, 43] and retrospective analyses [8]; however, there are currently few studies focusing specifically on CS3, the final course in most introductory sequences that generally covers some combination of data structures and algorithms. Current retention techniques involve *pipelining* students: introducing them to CS early and swiftly progressing them from one intro course to the next. The loss of certain populations (e.g., women) is typically described as a “leak” in the pipeline [10, 39].

With the increase of gatekeeping CS majors and courses to address the boom in undergraduate CS interest [13, 30], CS3 is sometimes the final prerequisite to declaring a CS-related major and thus accessing multiple upper-level courses, which increase students’ depth of knowledge. Retention efforts at the CS3 level are important for diversity, even if direct attrition rates are more similar [2]. Furthermore, in settings where major standing is required for enrollment in upper CS courses, major declarations can restrict access to resources and overall engagement. By CS3, students have shown an interest in Computer Science and built up programming experience and mathematical thinking. Thus, key factors found by previous studies at lower levels, where these properties generally do not hold, may not be applicable at higher levels. Additionally, populations that typically “leak from the pipeline” in CS1 and CS2 may no longer leave the major once they complete CS3, and techniques directed at them might best be leveraged for different populations more at risk at the CS3 level.

We develop a clearer understanding of major declarations and attrition rates at the CS3 level. We hypothesize that rates of students leaving CS, despite reporting an intention to stay (*attrition*) at the CS3 level correlate with biological sex, student grades, and student progression through the core CS curriculum (*fast-tracking*). We observe that both perceptions and outcomes may be factors when analyzing differences between major-declaring and non-declaring students in the same course. Thus, we develop a methodology to combine archival course data with university student record data. Our IRB-approved data set includes over 500 students enrolled in

the same semester of CS3 at a large public university. First, we conduct a partial replication study of previous findings [2] to see if sex relates to declaration rates. Next, we draw on previous findings on the correlation between grades and persistence through the major [21] to perform a similar study of grades and declaration rates. Then, we study the effects of a swift progression through the standardized portion of a CS curriculum, which we term *fast-tracking*. Finally, we study autograder usage, to determine whether there is a noticeable difference between retained students and those who leave the major, and between fast-tracked and non-fast-tracked students.

Our statistical analyses suggest that sex alone does not correlate with declaration rates, but grades and swift progression through the CS curriculum do. While we find the pipeline at CS3 to be leak-resistant for women, we identify a different population that is more at risk of leaving the major: students who take a longer gap between CS2 and CS3, or who take CS3 later in their college careers. These students show a 13 percentage point reduction in retention ($p = 0.00005$), indicating that these factors are very relevant for retention at the CS3 level. Additionally, our study of autograder usage and grades shows statistically significant differences across project submissions ($p = 0.00003$). Our findings suggest that interventions at the CS3 level should target students who arrive to the course later and focus on developing better autograder habits to increase grades.

To summarize, this work makes the following contributions:

- A combined analysis of archival CS3 course data and university-wide student records to study major declarations and attrition.
- A statistical analysis of fast-tracking (student progress through the core CS curriculum), and its effects at the CS3 level.
- A study of the relationship between student submissions to an autograder and major declarations.

2 BACKGROUND AND RELATED WORK

While attrition rates of men and women at the CS3 level have not been shown to be significantly different [2], to the best of our knowledge, no research specifically focuses on CS3 students who *explicitly intend to major in CS*. Attrition can be a particularly negative experience for these students, as they have already made a significant time investment, and changing paths could increase the amount of time needed to complete a degree, with financial impacts [25]. Studies of attrition at the CS1 level find no single contributing factor, instead identifying a combination of time, motivation, and comfort [22, 40, 42]. Our study produces similar findings at the CS3 level: we identify several factors contributing to attrition, but no one of these factors is sufficient for classifying students.

Research on women moving through the CS curriculum is often referred to as the *pipeline model*, with attrition being a *leak* in the pipeline [9, 41]. We are particularly interested in a facet of this pipeline we refer to as *fast-tracking*, wherein students progress quickly from one prerequisite course to the next. While students may benefit from taking courses in quick progression, because prior knowledge is still fresh [37], differences in declaration rates between fast-tracked students and non-fast-tracked students may indicate that a course should be adjusted to support students taking a more circuitous route to end up in CS, which is common among

women [35, 41], returning adult students [29, 45], and transfer students [38]. Furthermore, potential resources provided to these students might need to be adjusted, for example integrating additional tutoring for prerequisite concept review into a CS3 course, rather than offering another intermediary course before CS3.

Female students are not found to leave male-dominated fields (e.g., Computer Science) at higher rates than other fields [33]. Previous studies, however, indicate that student perceptions of their own effort and effort relative to their peers differ by gender, even when researchers correct for grades [23]. Higher grades are linked with higher persistence rates through the major, yet research shows that women often choose not to pursue Computer Science after introductory CS courses, despite having similar or higher grades than men [21], because their perception of effort is different. We test whether these findings are supported at the CS3 level, and what significant differences should be accounted for.

Identifying students at risk of leaving the major is necessary for instructional staff to intervene. Models for predicting student retention in CS have been proposed based on student effort, comfort level, instructional practices, degree usefulness, and cognitive gains [3, 17, 31, 40]. For example, Ahadi et al. trained a classifier on source code snapshots of student assignments to predict high- and low-performing students in CS1 [1] and Castro-Wunsch et al. used neural nets trained on similar small coding exercises [12]. These existing approaches can be challenging to implement at the CS3 level due—in part—to larger, more comprehensive projects and assignments. Automatic grading (autograders) have become increasingly common, especially in large departments [16, 28]. Studies show that automatically generated feedback (e.g., autograder reports) can positively impact student learning via timeliness and perceived constructiveness [5, 15, 19]. In this work, we investigate autograder usage and identify key differences between retained students and those who leave the major. We also discuss the applicability of these findings to efforts to identify at-risk students.

3 RESEARCH QUESTIONS

With this background in mind, we present the research questions guiding the remainder of this work. We first replicate findings from previous studies on major attrition. We also consider new aspects that we hypothesize impact students at the CS3 level, including autograder usage and fast-track progression of students through introductory CS courses. We discuss our data sources and results in sections 4 and 5, respectively.

The following four research questions guide our data analysis:

- RQ1:** Is there a correlation between biological sex and declaration rates for CS-related majors?
- RQ2:** Is there a correlation between student grades in CS3 and declaration rates of CS-related majors?
- RQ3:** How does student *fast-tracking* through core CS courses relate to declaration rates of CS-related majors?
- RQ4:** Can autograder data help identify CS3 students at risk of leaving CS-related majors?

4 DATA SOURCES

To answer our research questions, we combine archival course data with university-wide student records to allow for both qualitative

and quantitative analyses. In this section we describe both data sets, and we summarize the metrics we use as part of our analyses.

Our data collection protocol and FERPA considerations were approved by the University’s IRB (HUM00159578) in conjunction with two Memoranda of Understanding (MOU) with the Computer Science Department.

4.1 Archival Course Data

We collected archival course data from 2019 for one semester of CS3 (Data Structures and Algorithms) at the University of Michigan, a large Midwestern US public university. This course consisted of 3 contact hours a week of lecture led by professors, 2 contact hours of a discussion section led by a graduate or undergraduate instructional assistant, 10 lab assignments, 4 medium-sized programming projects, a midterm, and a final exam. The course covered data structures and introductory algorithms (e.g., stacks, queues, priority queues, trees, sorting, and graph algorithms), as well as time and space optimization. Students completed programming projects individually, but had the option of completing lab assignments with a partner. For all programming assignments (both labs and projects), feedback was given via an autograder. Our archival data consists of student lab answers, submissions to the autograder, grades and letter grade cutoffs.

We exclude from our data sets any students who chose to opt out of research, students who were younger than 18 years of age at the time of taking the course, and those without a final grade for the course (who we assume dropped the course or received an Incomplete). We infer students’ “intention to declare a CS-related major” from a voluntary survey component in a mid-semester lab. Students who did not complete this assignment are dropped from our data sets. In total, we removed 15/597 students from our data.

In the remainder of this subsection, we highlight the most relevant aspects of this archival data.

4.1.1 Lab Survey Assignment. We analyze student answers to a voluntary survey component of a lab completed at the half-way mark (just after the midterm). This survey covered demographics, intentions to declare a CS-related major, perceptions of effort in the course, and perceptions of effort relative to peers.

We are particularly interested in factors that impact a student’s decision to declare a CS major or follow through on an intent to declare a CS major. Students were asked whether they had already declared, or planned to declare, the CS major or minor. We group students who had declared a CS major or planned to declare a CS major as students who *explicitly intend to declare the major*. We focus on students who intend to declare a CS major, as opposed to a minor, because our university limits course enrollment for non-majors. Thus, the decision to declare a major entails more time with the program, as well as access to more of the program’s resources.

Additionally, using a Likert scale, students answered four questions rating the effort needed to achieve their desired grade and also the effort they felt expended relative to their peers.

4.1.2 Autograder Submissions. Our anonymized autograder data set includes the timestamp and score of each submission made by a student to the autograder for each project, as well as all of the

Table 1: Student CS Major Declarations by Intention to Declare and Sex

	Intend to Declare			Do Not Intend		
	M	F	Total	M	F	Total
Declare	273	79	352	70	27	97
Do Not Declare	30	9	39	67	27	94
Total	303	88	391	137	54	191

submitted files. From student autograder submissions, we calculate three additional metrics: average number of submits to the autograder per project, average time spent with the autograder per project, and maximum score per day the autograder is open for a project. The average time spent with the autograder per project is the sum of the number of days between a student’s first submission and the day the project was due, divided by the number of projects (4). The average number of submits is calculated similarly. The maximum score per day is a rolling cumulative maximum of the best submission for a given day and any previous day. Students who have not yet submitted have a maximum score of 0.

4.1.3 Grades. Our data set includes grades for each student, including individual lab 1–10 scores, individual project 1–4 scores, midterm exams, final exams, averages of each section (Labs, Projects, Exams), and total course score. Grade cutoffs are provided to bin total course scores into letter grades ranging from F to A+.

4.2 University Student Records

Our study also uses university-level student records up to and through the semester *following* the course we analyze. At the University we are studying, most students declare a major during the semester they are enrolled in CS3 or the subsequent semester, because major or minor status is required to enroll in upper level CS courses, and thus we are confident we have captured the major declarations for most students in our study.

Within the data set, we look at variables for sex,¹ major declarations, and number of semesters enrolled in the university. To account for multiple computing-related majors, we consider two Computer Science degrees, one Computer Engineering degree, and two Data Science degrees to all be “majoring in Computer Science”.

5 RESULTS

In this section we analyze the data sources presented in Section 4 to test the hypotheses that biological sex, grades, fast-tracking, and autograder usage are correlated with major attrition rates at the CS3 level. We address the research questions presented in Section 3, requiring the standard statistical significance level of $p < 0.05$.

5.1 Sex and Major Declarations (RQ1)

We begin by replicating the study of Baer and DeOrio [2], who found that there is no significant difference between the proportion of women and men students who are retained at the CS3 level.

¹While student gender might be a more precise and relevant metric for our study, the data available to us only contained biological sex.

As shown in Table 1, of the 142 female students in our data set, 106 went on to declare a CS major (74.65%), while 343 of the 440 male students went on to declare a CS major (77.95%). Using a 2 proportion z-test, we find no significant difference in the proportion of female and male students in our data set who declare a CS major ($p = 0.42$), supporting the findings of Baer and DeOrio.

We extend this analysis to also consider a student’s intention to declare a CS-related major. In particular, we test if female students who intend to declare a CS major end up not declaring at *different* rates from their male counterparts. A 2 proportion z-test with Yates Continuity correction to account for fewer than 10 female students in this category does not yield statistically significant results ($p = 0.5$). That is, female students are no more likely than male students to not declare a CS major despite having the intention to. However, we note that female students still report significantly lower senses of belonging and higher expected effort in CS3 ($p = 0.00002$).

From these findings, we deduce that the pipeline model at the CS3-major junction seems to be working as effectively for men as for women: sex is not a relevant variable when predicting a student majoring after completing CS3, and candidate interventions related to gender are likely more effective in CS1 and CS2. However, as we will demonstrate, this does not mean the pipeline is fully successful. Rather, interventions in CS3 should target a different population for whom the pipeline is less effective: non-fast-tracked students.

Confirming previous results, sex alone does not predict major attrition for CS3 ($p = 0.42$) in our data. Further, there is no difference when considering intention to declare ($p = 0.5$). We must consider alternate factors or populations for CS3.

5.2 CS3 Grades and Major Declarations (RQ2)

Having established that sex does not play a significant role in attrition rates at the CS3 level, we next consider student grades, which have been found to correlate with persistence in the major [21]. We leverage these results to aid our study of student progression through the introductory course sequence in Section 5.3.

First, we study the raw total course scores for students who declare the major and those who do not. As shown in Table 1, of 582 students, 449 went on to declare a CS major. The mean total class score of all students is 80.02 points, but is 80.75 for those who declared and 77.55 for those who did not. Those differences in means are statistically significant ($p = 0.02$ via a Mann-Whitney U test for similar non-normal distributions, Shapiro-Wilk $p = 1.238 \times 10^{-15}$, $p = 8.823 \times 10^{-11}$). A difference in grades between those who declare and those who do not is perhaps expected. We present these results as a baseline for comparison with previous findings.

For major retention, we focus on students who express intention to declare a CS major. We repeat this analysis on those students who intend to declare a CS major ($n = 391$), splitting on major declaration ($n = 352$ declare vs. $n = 39$ leave). We find that the mean score for those who declare is 81.03 points, a much higher value than the mean score for those who *do not* declare, 70.05 points. While the two populations are independent, their distributions are not similar enough to test for a difference in means. However, the almost 11 point difference in grades between the two populations




Figure 1: Major declaration rates of students who intend to major in CS, grouped by semester since they took CS2. Note that only semesters with more than one student are shown.

correlates with significantly different letter grades (low C vs. high B-). While correlation is not causation, this concretely calls out grade performance in CS3 as a factor in CS3-to-major retention. Next, we analyze one potential factor—fast-tracking of students through core courses—and its relationship with grades.

Students who declare a CS major have a higher mean score in CS3 than those who do not ($p = 0.02$). Critically, students who *intend* to declare a major but ultimately leave perform a letter grade lower (on average) than students who declare.

5.3 Fast-Tracking Major Declarations (RQ3)

Research on CS retention (especially with regard to women in CS) often mentions the *pipeline model* [10, 39]. Students quickly flow from one CS class to the next, starting with pre-college preparation courses and ending with a CS-related career after graduation. In the college portion of this pipeline, there is additional pressure for students to take courses in rapid succession (i.e., without a semester “break” from CS) early in their college careers, which we term *fast-tracking*. We investigate how fast-tracking relates to major declarations. As such, we restrict our analysis in this subsection to only those students reporting an intention to declare a major.

Figure 1 shows the percentage of students who declare a CS major when completing CS3 m semesters after CS2. For example, 93% of students taking CS3 immediately after CS2 ($m = 1$) declare the major, while 7% do not declare. For visual clarity, we omit semesters with only one student; all students are considered in our statistical tests. An increasing percentage of students do not declare as this gap increases. The mean number of semesters since taking CS2 is significantly lower for students who declare a CS major (1.54 semesters vs 2.36 for no declaration, Mann-Whitney U Test $p = 0.00007$). This is evidence that those who do not fast-track the introductory courses are less likely to declare a CS major.

We previously found that grades can help predict whether a student will declare a major (Section 5.2). It is therefore possible that students who fast-track CS courses have higher grades and thus end up declaring the major. To test for this, we compare the mean grade of students who take CS3 immediately or take a one-semester break after CS2 (80.48%) and for students who take a two-semester or more break after CS2 (69.22%). This eleven point




Figure 2: Mean max score for Project 1 per open autograder day for students who intend to declare a CS major and go on to declare (“stay”) and students who do not (“leave”).

difference in mean grade is statistically significant (Mann-Whitney U Test, $p = 0.00008$). Therefore, students who fast-track the introductory sequence achieve higher grades and are more likely to follow through with major declaration.

The pipeline model for CS also focuses on early introduction and recruitment, including for underrepresented students [20, 44]. To analyze the effect of early recruitment on retention, we compare the mean months of college before completing CS3 for those who do declare (25.26 months, $n = 352$) to those who do not (30.67 months, $n = 39$). Students who declare the major complete CS3, on average, in the first two years of study vs. three years for those who do not (Mann-Whitney U test, $p = 0.00006$).

Using these findings, we compare the proportion of fast-tracked students (those who complete CS3 in the first two years with no more than a one-semester gap with CS2) declaring CS majors to those who are not fast-tracked. We find that 13 percentage point more fast-tracked students declare the major (93.84% vs. 80.87%, one-sided 2 proportion z-test, $p = 0.00005$). Note that this result only considers students who express interest in declaring the major.

Students who declare a CS major fast-track introductory CS courses ($p = 0.00007$), have higher course grades ($p = 0.00008$), and complete CS3 earlier in college ($p = 0.00006$). Fast-tracked students are 13 percentage points more likely to follow through with declaring a CS major.

5.4 Autograders and Major Declarations (RQ4)

Finally, we consider the relationship between autograder interactions and major declarations. Autograders are increasingly common in computing classes but are relatively understudied in terms of retention. Given that grades and fast-tracking are both relevant to CS3 major declarations, we hypothesize autograder interactions to be an additional source of relevant information (e.g., these systems directly influence grades, and some non-fast-tracked students, such as non-traditional students, may have less preparation involving their use).

We consider autograder interactions both in terms of submission counts and number of days of use (Section 4.1.2). Students who go on to declare a major use the autograder for an average of




Figure 3: Mean max score for Project 1 per open autograder day for fast-tracked and non-fast-tracked students who intend to declare a CS major.

7.63 days per project, but those who do not declare only use 6.06 days per project. This difference is statistically significant (Mann-Whitney U test, $p = 0.0083$). We observe a similarly significant difference in the number of submissions to the autograder (16.2 vs. 12.5 submissions, Mann-Whitney U test, $p = 0.0005$). Further, using linear regression, we identify a positive relationship between number of submissions and duration of usage ($p = 2.2 \times 10^{-16}$, $R^2 = 0.65$). Thus, students who submit early and often are more likely to declare a CS major than those who do not, highlighting a potential for particular interventions (e.g., focusing on improving autograder habits for CS3 non-fast-tracked students).

Further, we observe a significant grade disparity between students who declare a major and those who do not. Figure 2 shows the average maximum score over time on the first project. A Wilcoxon Signed Rank Test for the pairwise difference of means finds the trend for those who stay to be significantly higher than for those who leave ($p = 0.00003$). Similarly, Figure 3 shows that students who are fast-tracked have a higher project average throughout than those who are not ($p = 0.00003$). This same trend exists for the other projects in our data set.

Students who intend to declare a CS major and go on to declare spend more time with the autograder ($p = 0.0083$) and use more submits ($p = 0.0005$) than those who intend to declare a CS major but do not do so. While preliminary, these results suggest that autograder interactions may be a fruitful source of retention information.

6 DISCUSSION

By the time students are enrolled in CS3, they have already invested significantly in their CS education. Retention efforts at the CS3 level benefit from the observation that these students typically *want* to study Computer Science. CS3 is a critical transition point in the CS “pipeline model” when students declare the major. Our results suggest that this stage of the pipeline is generally resilient to “leaks” based on biological sex. However, higher grades are associated with higher retention rates of CS majors, and the current curriculum significantly favors fast-tracked students. These results motivate additional investigations of the impact of fast-tracking

and disparities in student experience, including focusing on courses earlier in the CS core curriculum, with the goal of identifying how these factors fluctuate as students progress.

The indication that fast-tracking can help retain students suggests we should make attempts to introduce students to CS earlier and support taking CS2 and CS3 back-to-back. This can be a particular challenge for groups of students (e.g., women, returning adult students, and transfer students) who are less likely to be on a fast-track through the curriculum or are target demographics for longer on-boarding programs [29, 35, 38, 41, 45]. In addition to encouraging students to start CS early, this suggests that we should also provide additional support and resources for students who discover CS later.

Our results indicate that there is a nontrivial number of students in CS3 (11%) who *want* to declare a major, but despite their existing time investment *do not*. Helping instructors identify these students would allow for interventions that could retain these students. Our results indicate specific targeting of female students may be less appropriate at the CS3 level than in earlier courses. Instead, our results suggest that identifying non-fast-tracked students may be a more fruitful approach for intervention efforts seeking to increase major retention. While grades do appear to be correlated with students leaving the major, this information is counter-intuitively difficult to apply to classify students leaving the major. Our initial efforts to use both ensemble- and time-series-based classifiers produced results that were not significantly better than targeting a known focus group, even when training classifiers separately on fast-tracked and non-fast-tracked students. This is likely due to the wide variance in grades for students who do remain in the major (students who declare the major can often have lower grades than “mask” the at-risk students). As such, using grades alone to identify students for a candidate intervention seems to be insufficient. Instead, there appears to be a much more complex interaction between grades and aspects of student’s identity and experience that result in a decision to leave the major, similar to the multifaceted reasons for dropping CS1 found by Kinnunen and Malmi [22]. Though we know that grades are correlated with major declarations (RQ2) and students who leave the major underperform on the autograder (RQ4), determining an appropriate combination of metrics to identify individual students for interventions warrants further investigation.

Our archival data also indicates that 16% of students in the course *do not* intend to declare a major, but ultimately *do*. While analysis of this cohort falls outside the scope of this investigation, this phenomenon is worth studying. Factors and interventions that attract initially disinterested students may help retain a broader student body. Such findings also have the potential to help departments retain students who are interested in declaring the major but ultimately leave the program.

7 THREATS TO VALIDITY

In this section we discuss two key threats to validity and the mitigation techniques we used when designing our study.

Our assignment data came from one semester of CS3 and our findings may not generalize to other universities or other semesters. We mitigate this threat by both considering a large, multi-section course (over 500 students) and also using other data sources (e.g.,

university records data) that span multiple semesters. Further, we first replicate results of previous studies with our data to gain confidence that we are not studying an anomalous semester. We leave a multi-institutional study for future work.

Our study may also admit threats to internal validity. We mitigate such threats by rigorously sanitizing our data (including removing incomplete data that might confound our results) and cross-checking metrics across our multiple data sources wherever possible. Further, we conducted our study on archival data meaning that students’ outcomes were not biased by any active interventions or study.

8 CONCLUSIONS

Retaining a diverse student body of CS majors remains an open challenge for our field. In this work, we extend previous studies that focus on student retention in CS1 and CS2 to include analysis of CS3 students. We analyze archival data and university student records for over 500 students enrolled in CS3 at the University of Michigan, a large Midwestern public university in the US.

First, we replicate existing findings that biological sex alone does not impact major attrition and the current pipeline is leak-resistant at the CS3 level. Next, we study the relationship between grades and major declaration, finding that CS3 students who declare a CS major outperform those who do not declare by a letter grade, on average. Third, we find students who take CS2 and CS3 back-to-back (and within the first two years) are more likely to declare a CS major than students who are not fast-tracked. Finally, we observe statistically significant differences in autograder usage between students who major and students who leave the program, as well as fast-tracked and non-fast-tracked students.

Our findings demonstrate that fast-tracking of students through core CS courses helps retain students in the major, but fails to adequately support those students who reach CS on a longer journey. We expose new avenues of inquiry for identifying best practices to retain students who have reached CS3, including study of interventions at the CS3 level that both target non-fast-tracked students and also focus on developing better autograder habits.

ACKNOWLEDGMENTS

This work was funded, in part, by the NSF (1629450, 1763674, 1908633) and the Rackham Graduate School at the University of Michigan. We also thank Jonathan Dorn, Matt Higham, Zohreh Shirafi, and Karthik Urs for their guidance on various statistical techniques used in this work.

REFERENCES

- [1] Alireza Ahadi, Raymond Lister, Heikki Haapala, and Arto Vihavainen. 2015. Exploring Machine Learning Methods to Automatically Identify Students in Need of Assistance. In *Proceedings of the Eleventh Annual International Conference on International Computing Education Research* (Omaha, Nebraska, USA) (ICER '15). Association for Computing Machinery, New York, NY, USA, 121–130.
- [2] Amy Baer and Andrew DeOrio. 2020. A Longitudinal View of Gender Balance in a Large Computer Science Program. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*. ACM, Portland OR USA, 23–29.
- [3] L. Barker, C. L. Hovey, and L. D. Thompson. 2014. Results of a large-scale, multi-institutional study of undergraduate retention in computing. In *2014 IEEE Frontiers in Education Conference (FIE) Proceedings*. 1–8.
- [4] Lecia J. Barker, Charlie McDowell, and Kimberly Kalahar. 2009. Exploring factors that influence computer science introductory course students to persist in the major. *ACM SIGCSE Bulletin* 41, 1 (March 2009), 153–157.

- [5] Leopold Bayerlein. 2014. Students' feedback preferences: how do students react to timely and automatically generated assessment feedback? *Assessment & Evaluation in Higher Education* 39, 8 (2014), 916–931.
- [6] Sylvia Beyer. 2014. Why are women underrepresented in Computer Science? Gender differences in stereotypes, self-efficacy, values, and interests and predictors of future CS course-taking and grades. *Computer Science Education* 24, 2-3 (July 2014), 153–192. Publisher: Routledge.
- [7] Sylvia Beyer, Kristina Rynes, Julie Perrault, Kelly Hay, and Susan Haller. 2003. Gender differences in computer science students. In *Proceedings of the 34th SIGCSE technical symposium on Computer science education - SIGCSE '03*. ACM Press, Reno, Nevada, USA, 49.
- [8] Maureen Biggers, Anne Brauer, and Tuba Yilmaz. 2008. Student Perceptions of Computer Science: A Retention Study Comparing Graduating Seniors with Cs Leavers. In *Proceedings of the 39th SIGCSE Technical Symposium on Computer Science Education* (Portland, OR, USA) (SIGCSE '08). Association for Computing Machinery, New York, NY, USA, 402–406.
- [9] Jacob Clark Blickenstaff*. 2005. Women and science careers: leaky pipeline or gender filter? *Gender and Education* 17, 4 (2005), 369–386.
- [10] Tracy Camp. 2002. The incredible shrinking pipeline. *ACM SIGCSE Bulletin* 34, 2 (2002), 129–134.
- [11] Lori Carter. 2006. Why Students with an Apparent Aptitude for Computer Science Don't Choose to Major in Computer Science. In *Proceedings of the 37th SIGCSE Technical Symposium on Computer Science Education* (Houston, Texas, USA) (SIGCSE '06). Association for Computing Machinery, New York, NY, USA, 27–31.
- [12] Karo Castro-Wunsch, Alireza Ahadi, and Andrew Petersen. 2017. Evaluating Neural Networks as a Method for Identifying Students in Need of Assistance. In *Proceedings of the 2017 ACM SIGCSE Technical Symposium on Computer Science Education* (Seattle, Washington, USA) (SIGCSE '17). Association for Computing Machinery, New York, NY, USA, 111–116.
- [13] Committee on the Growth of Computer Science Undergraduate Enrollments, Board on Higher Education and Workforce, Policy and Global Affairs, Computer Science and Telecommunications Board, Division on Engineering and Physical Sciences, and National Academies of Sciences, Engineering, and Medicine. 2018. *Assessing and Responding to the Growth of Computer Science Undergraduate Enrollments*. National Academies Press, Washington, D.C. Pages: 24926.
- [14] Computing Research Association. 2017. *Generation CS: Computer Science Undergraduate Enrollments Surge Since 2006*. Technical Report.
- [15] Phillip Dawson, Michael Henderson, Paige Mahoney, Michael Phillips, Tracii Ryan, David Boud, and Elizabeth Molloy. 2019. What makes for effective feedback: staff and student perspectives. *Assessment & Evaluation in Higher Education* 44, 1 (Jan 2019), 25–36.
- [16] Christopher Douce, David Livingstone, and James Orwell. 2005. Automatic Test-Based Assessment of Programming: A Review. *J. Educ. Resour. Comput.* 5, 3 (Sept. 2005), 4–es.
- [17] Michail N. Giannakos, Ilias O. Pappas, Letizia Jaccheri, and Demetrios G. Sampson. 2017. Understanding student retention in computer science education: The role of environment, gains, barriers and usefulness. *Education and Information Technologies* 22, 5 (Sep 2017), 2365–2382.
- [18] Graciela Gonzalez. 2006. A Systematic Approach to Active and Cooperative Learning in CS1 and Its Effects on CS2. *SIGCSE Bull.* 38, 1 (March 2006), 133–137.
- [19] J. W. Harrold and J. Sandland. 2018. The Influence of Immediate Homework Feedback on Student Performance and Satisfaction in an Engineering MOOC. In *2018 Learning With MOOCs (LWMOOCs)*. 90–93.
- [20] Andrea Jepson and Teri Perl. 2002. Priming the Pipeline. *SIGCSE Bull.* 34, 2 (June 2002), 36–39.
- [21] Sandra Katz, David Allbritton, John Aronis, Christine Wilson, and Mary Lou Soffa. 2006. Gender, achievement, and persistence in an undergraduate computer science program. *ACM SIGMIS Database: the DATABASE for Advances in Information Systems* 37, 4 (Nov. 2006), 42–57.
- [22] Päivi Kinnunen and Lauri Malmi. 2006. Why students drop out CS1 course?. In *Proceedings of the 2006 international workshop on Computing education research - ICER '06*. ACM Press, 97.
- [23] H. Bruce Lammers, Tina Kiesler, Mary T. Curren, Deborah Cours, and Brian Connett. 2005. How Hard Do I Have to Work? Student and Faculty Expectations Regarding University Work. *Journal of Education for Business* 80, 4 (2005), 210–213.
- [24] Karyn L. Lewis, Jane G. Stout, Noah D. Finkelstein, Steven J. Pollock, Akira Miyake, Geoff L. Cohen, and Tiffany A. Ito. 2017. Fitting in to Move Forward: Belonging, Gender, and Persistence in the Physical Sciences, Technology, Engineering, and Mathematics (pSTEM). *Psychology of Women Quarterly* 41, 4 (Dec 2017), 420–436.
- [25] Jon Marcus. 2018. Switching majors is adding time and tuition to the already high cost of college. *Hechinger Report* (Dec 2018).
- [26] Beryl Nelson. 2014. The data on diversity. *Commun. ACM* 57, 11 (Oct 2014), 86–95.
- [27] An Nguyen and Colleen M. Lewis. 2020. Competitive Enrollment Policies in Computing Departments Negatively Predict First-Year Students' Sense of Belonging, Self-Efficacy, and Perception of Department. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*. ACM, Portland OR USA, 685–691.
- [28] Pete Nordquist. 2007. Providing Accurate and Timely Feedback by Automatically Grading Student Programming Labs. *J. Comput. Sci. Coll.* 23, 2 (Dec. 2007), 16–23.
- [29] Daphne W. Ntiri. 2001. Access to Higher Education for Nontraditional Students and Minorities in a Technology-Focused Society. *Urban Education* 36, 1 (2001), 129–144.
- [30] E. Patitsas, M. Craig, and S. Easterbrook. 2016. How CS departments are managing the enrolment boom: Troubling implications for diversity. In *2016 Research on Equity and Sustained Participation in Engineering, Computing, and Technology (RESPECT)*. 1–2.
- [31] Markeya S. Peteranetz and Leen-Kiat Soh. 2020. A Multi-Level Analysis of the Relationship between Instructional Practices and Retention in Computer Science. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education* (Portland, OR, USA) (SIGCSE '20). Association for Computing Machinery, New York, NY, USA, 37–43.
- [32] J. Raigoza. 2017. A study of students' progress through introductory computer science programming courses. In *2017 IEEE Frontiers in Education Conference (FIE)*. 1–7.
- [33] Catherine Riegle-Crumb, Barbara King, and Chelsea Moore. 2016. Do They Stay or Do They Go? The Switching Decisions of Individuals Who Enter Gender Atypical College Majors. *Sex Roles* 74, 9–10 (May 2016), 436–449.
- [34] Linda J. Sax, Jennifer M. Blaney, Kathleen J. Lehman, Sarah L. Rodriguez, Kari L. George, and Christina Zavala. 2018. Sense of Belonging in Computing: The Role of Introductory Courses for Women and Underrepresented Minority Students. *Social Sciences* 7, 8 (Jul 2018), 122.
- [35] Linda J. Sax, Kathleen J. Lehman, and Christina Zavala. 2017. Examining the Enrollment Growth: Non-CS Majors in CS1 Courses. In *Proceedings of the 2017 ACM SIGCSE Technical Symposium on Computer Science Education*. ACM, 513–518.
- [36] Duane F. Shell, Leen-Kiat Soh, Abraham E. Flanigan, and Markeya S. Peteranetz. 2016. Students' Initial Course Motivation and Their Achievement and Retention in College CS1 Courses. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education* (Memphis, Tennessee, USA) (SIGCSE '16). Association for Computing Machinery, New York, NY, USA, 639–644.
- [37] Matthew F. Tennyson and Marc Beck. 2018. A Study of Knowledge Retention in Introductory Programming Courses. *J. Comput. Sci. Coll.* 33, 4 (April 2018), 13–20.
- [38] Barbara K Townsend. 2000. Transfer students' institutional attendance patterns: A case study. *College and University* 76, 1 (2000), 21.
- [39] R. Varma and H. Hahn. 2008. Gender and the pipeline metaphor in computing. *European Journal of Engineering Education* 33, 1 (2008), 3–11.
- [40] Philip R. Ventura. 2005. Identifying predictors of success for an objects-first CS1. *Computer Science Education* 15, 3 (Sep 2005), 223–243.
- [41] Anna Vitores and Adriana Gil-Juárez. 2016. The trouble with 'women in computing': a critical examination of the deployment of research on the gender gap in computer science. *Journal of Gender Studies* 25, 6 (Nov 2016), 666–680.
- [42] Brenda Cantwell Wilson. 2002. A Study of Factors Promoting Success in Computer Science Including Gender Differences. *Computer Science Education* 12, 1-2 (March 2002), 141–164. Publisher: Routledge.
- [43] Dana Wortman and Penny Rheingans. 2007. Visualizing Trends in Student Performance across Computer Science Courses. *SIGCSE Bull.* 39, 1 (March 2007), 430–434.
- [44] Ken Yasuhara. 2005. Choosing computer science: Women at the start of the undergraduate pipeline. In *Proceedings of the American Society for Engineering Education Annual Conference*.
- [45] Anna Zajacova, Scott M. Lynch, and Thomas J. Espenshade. 2005. Self-Efficacy, Stress, and Academic Success in College. *Research in Higher Education* 46, 6 (Sep 2005), 677–706.